arXiv:1410.7852v1 [cs.LG] 29 Oct 2014

A Markov Decision Process Analysis of the Cold Start Problem in Bayesian Information Filtering

Xiaoting Zhao, Peter I. Frazier
School of Operations Research and Information Engineering
Cornell University, Ithaca, NY 14850
xzhao@cornell.edu, pf98@cornell.edu

October 30, 2014

Abstract

We consider the information filtering problem, in which we face a stream of items, and must decide which ones to forward to a user to maximize the number of relevant items shown, minus a penalty for each irrelevant item shown. Forwarding decisions are made separately in a personalized way for each user. We focus on the cold-start setting for this problem, in which we have limited historical data on the user's preferences, and must rely on feedback from forwarded articles to learn which the fraction of items relevant to the user in each of several item categories. Performing well in this setting requires trading exploration vs. exploitation, forwarding items that are likely to be irrelevant, to allow learning that will improve later performance. In a Bayesian setting, and using Markov decision processes, we show how the Bayes-optimal forwarding algorithm can be computed efficiently when the user will examine each forwarded article, and how an upper bound on the Bayes-optimal procedure and a heuristic index policy can be obtained for the setting when the user will examine only a limited number of forwarded items. We present results from simulation experiments using parameters estimated using historical data from arXiv.org.

1 Introduction

We study the information filtering problem, in which a user faces a stream of time-sensitive items (emails, blog posts, scientific articles), some of which are interesting to the user, but many of which are uninteresting. We wish to design an automatic system that automatically filters this stream, showing as many relevant items to the user as possible, while showing few irrelevant items.

When historical relevance data from the user is abundant, we can train a statistical classifier, and forward only those items predicted to be relevant. However, when historical data from the user is limited, e.g., because the user is new, or because we are dealing with new kinds of items, we face the *cold start problem*, in which we do not have enough training data to build a reliable classifier.

In this setting, it may be advantageous to *explore*, i.e., forward some items predicted to be irrelevant, just to learn their true relevance and improve future predictions. Too much exploration, however, will lead to forwarding many irrelevant items. Thus, an information filtering system should also put some weight on *exploitation*, i.e., forwarding only those items predicted to be relevant.

We study this tradeoff between exploration vs. exploitation in a Bayesian setting, using a Markov decision process (MDP). We show how the MDP defining the Bayes-optimal algorithm for making forwarding decisions may be solved efficiently using a decomposition, when irrelevant items are penalized by a user-specified cost per item shown. We then show how to use this solution to provide a ranking over items when a cost-per-item is unknown not given and users are only willing to examine a limited number of items, but users examine each forwarded item.

Exploration vs. exploitation has been studied extensively in the context of the multi-armed bandit problem in both Bayesian treatments (Gittins and Jones 1974, Whittle 1980, Gittins et al. 2011), and non-Bayesian treatments (Auer et al. 1995, 2002). This tradeoff between exploration and exploitation, which appears in other problem domains including reinforcement learning (Kaelbling et al. 1998, Sutton and Barto 1998, Jaksch et al. 2010), approximate dynamic programming (Powell et al. 2004, Powell 2007), revenue management (Araman and Caldenty 2009, Besbes and Zeevi 2009, den Boer and Zwart 2013), optimization algorithms (Xie and Frazier 2013a, Frazier et al. 2008), and inventory control (Lariviere and Porteus 1999, Ding et al. 2002). In information retrieval problems, exploration vs. exploitation has also been studied in Zhang et al. (2003), Agarwal et al. (2009), Yue et al. (2009), Hofmann et al. (2013).

We are motivated by a personalized information system we are building for the electronic repository of scientific articles, arXiv.org. In popular categories like astro-ph (Astrophysics) and hep-th (high-energy physics), roughly 80 new articles are submitted each week (arXiv.org 2014), which creates a challenge for scientists who wish to remain abreast of new arXiv articles directly relevant to their research. Our experimental results use parameter settings estimated from historical data from arXiv.org.

This paper builds on the previous work Zhao and Frazier (2014), which considers an information filtering problem in a Bayesian setting, and uses dynamic programming to find the Bayes-optimal strategy for trading exploration and exploitation. There are two main differences between the model considered in that paper, and the one considered in the current paper. First, in Zhao and Frazier (2014), users provide *immediate* feedback on forwarded items, while in the current paper, we allow items to queue in the system until the next user visit. It is only upon visiting the system that the user provides feedback. This "periodic review" assumption is more realistic in many information filtering systems.Second, Zhao and Frazier (2014) assumes that users provide a unit cost for forwarding to the information filtering system, while in the current paper we provide a method for ranking results that allows this cost to be unknown.

Our focus on a Bayesian setting and Bayes-optimal procedures (rather than procedures that are just optimal up to a constant), is in contrast with the portion of the literature on multi-armed bandits that examines regret in a worst-case setting, and provides algorithms that have optimal

dependence on time or other problem parameters, but ignore constants. This focus allows us to apply our method profitably in small-sample regimes, where the best worst-case guarantee would be much worse than the best average-case guarantee, and where constants are just as important as the dependence on time. A downside of our focus on the Bayesian setting is that it requires us to choose a prior distribution to use when measuring performance. However, in many applied settings, including arXiv.org, we argue that a reasonable prior distribution can be estimated from historical data.

In Section 2, we formulate the information filtering problem with periodic reviews and a fixed unit cost for forwarding an item. In Section 3, we consider the case where the unit cost is unknown and there is a budget constraint on the total number of items that a user can view. We then show how to derive a ranking from this budgeted problem. Lastly, we show experimental results in Section 4.

2 Mathematical Model with a Unit Cost for Forwarding

We assume that items arrive to the system according to a Poisson process with rate $\lambda > 0$. Each item is categorized into (exactly) one of k categories, $\{1, ..., k\}$, and the category is observed as it enters the system. For systems without explicit categorization, the categories could be obtained by running a clustering algorithm on previous collected items in a pre-processing step. Let X_i denote

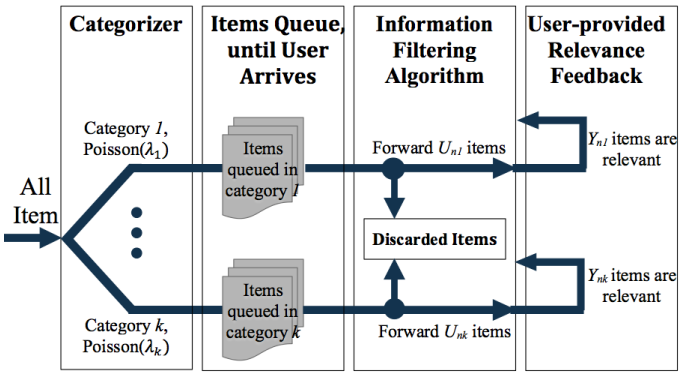

Figure 1: Schematic of the information filtering problem with periodic reviews.

the category of the i^{th} arriving item. We assume that the X_i are independent and identically distributed, and we let $p_x = P(X_i = x) > 0$. Thus, items in each category $x \in \{1, ..., k\}$ arrive according to a Poisson process with rate $\lambda_x = p_x \lambda$.

Each category x has some latent unobserved value $\theta_x \in [0, 1]$ measuring the probability that an item from category x is relevant to the user. Let $\theta = [\theta_1, ..., \theta_k]$. We place a Bayesian prior distribution on this θ, given by $\theta_x \sim \text{Beta}(\alpha_{0x}, \beta_{0x})$, with independence across x, for some parameters α_{0x} and β_{0x}, typically estimated using historical data from (other) long-time users on older items. In our model, θ is assumed to stay static over the user's lifetime.

The optimal tradeoff of exploration vs. exploitation will depend on how long the user interacts with our stream of items. Let T be the length of time that the user uses our information filtering system. T is unknown a priori, and we model it as an exponential random variable with parameter r.

The previous model in Zhao and Frazier (2014) provides a Bayes-optimal algorithm that analyzes the situation in which the user is always available to provide immediate feedback on each arriving item. However, in many real systems, users do not behave like this. Instead, they arrive periodically to review items that have queued in the system since their last visit.

Our model assumes that the user visits the system at time points separated by exponentially distributed inter-arrival times, which are independent and have parameter s. Let N be the number of user visits before T. Here, we assume that at each visit, he or she examines all items forwarded from the stream since the last visit. Below, we study a problem variant in which the number of items the user is willing to examine on each visit is constrained.

At the n^{th} user visit, the posterior on θ_x is Beta$(\alpha_{nx}, \beta_{nx})$, for some α_{nx}, β_{nx}. Based on $(\alpha_{nx}, \beta_{nx})$, we choose U_{nx}, denoting the maximum number of items to forward to the user from category x. For simplicity, we choose U_{nx} before observing L_{nx}, which is the number of items queued in category x, then we show $Z_{nx} = \min(L_{nx}, U_{nx})$ items from category x to the user in his or her n^{th} visit. For computationally convenience, we require $U_{nx} \leq M$, where $M < \infty$.

With the decision U_{nx} and forwarded Z_{nx} items, the user provides explicit feedback, denoted by Y_{nx}, reflecting the actual relevance of the shown items to the user. Conditioning on U_{nx}, θ_x and L_{nx}, Y_{xn} is binomial with a probability of θ_x being relevant to the user, that is,

$$Y_{nx} \mid \theta_x, U_{nx}, L_{nx} \sim \text{Binomial}\left(Z_{nx}, \theta_x\right).$$

We also assume there is a unit cost, c, for forwarding each item to the user. Thus, we collect reward $Y_{nx} - cZ_{nx}$ in each step. We define a policy π as a sequence of functions, $(\pi_1, \pi_2, ..., \pi_N)$, where each $\pi_n : (\{0, 1..., M\}^k \times \mathbb{N}^k \times \{0, 1..., M\}^k)^{n-1} \mapsto \{0, 1, ..., M\}^k$ maps history, $\{U_{\ell x}, L_{\ell x}, Y_{\ell x} : \ell \leq n-1, x \in \{1, ..., k\}\}$ into actions. Let Π be the set of all such policies. Our objective becomes to find an optimal policy $\pi \in \Pi$ that maximizes total expected reward:

$$\sup_{\pi \in \Pi} \mathbb{E}^{\pi} \left[\sum_{n=1}^{N} \sum_{x=1}^{k} (Y_{nx} - cZ_{nx}) \right]. \tag{1}$$

2.1 Solution and Computation Method

Due to the independence assumption across θ_x, we can decompose the original problem with k-categories into a sum of k independent sub-problems, each of which can be solved via stochastic dynamic programming. Equation (1) is rewritten as,

$$\sup_{\pi \in \Pi} E^{\pi} \left[\sum_{x=1}^{k} \sum_{n=1}^{N} (Y_{nx} - cZ_{nx}) \right] = \sum_{x=1}^{k} \sup_{\pi(x) \in \Pi(x)} E^{\pi(x)} \left[\sum_{n=1}^{N} (Y_{nx} - cZ_{nx}) \right],$$

Algorithm 1 Computation of $V_x^L(\alpha,\beta;\tilde{T})$ and $V_x^U(\alpha,\beta;\tilde{T})$

Require: γ, α_{0x}, β_{0x}, ξ_x, c, and N
  for $i = 0, ..., \tilde{T} + M$ **do**
    for $j = \max\{0, \tilde{T} + 1 - \alpha\}, ..., \tilde{T} + M - \alpha$ **do**
      Let $\alpha = \alpha_{0x} + i$, and $\beta = \beta_{0x} + j$.
      Let $V_x^L(\alpha,\beta;\tilde{T}) = \frac{E[\min(M,L_{1x})]}{(1-\gamma)(1-\gamma\xi_x)} \max\left\{0, \frac{\alpha}{\alpha+\beta} - c\right\}$ and $V^U(\alpha,\beta;\tilde{T}) = \frac{E[\min(M,L_{1x})]}{(1-\gamma)(1-\gamma\xi_x)}$.
    end for
  end for
  for $i = N_{ttl}, ..., 0$ **do**
    for $j = N_{ttl} - \alpha, ..., 0$ **do**
      Let $\alpha = \alpha_{0x} + i$, $\beta = \beta_{0x} + j$, and $\mu = \frac{\alpha}{\alpha+\beta}$.
      Let $V_x^L(\alpha,\beta;\tilde{T}) = \max_{0\le u\le M}\left\{(\mu - c)E(\min(u, L_{1x})] + \gamma E[V_x^L(\alpha_{1x},\beta_{1x}; N_{ttl})|U_{1x} = u]\right\}$,
      $V_x^U(\alpha,\beta;\tilde{T}) = \max_{0\le u\le M}\left\{(\mu - c)E(\min(u, L_{1x})] + \gamma E[V_x^U(\alpha_{1x},\beta_{1x}; N_{ttl})|U_{1x} = u]\right\}$,
    end for
  end for

where a policy $\pi(x)$ is a sequence of functions, $(\pi_1(x), \pi_2(x), ..., \pi_N(x))$, associated with category x and $\Pi(x)$ is the set of all $\pi(x)$. Each $\pi_n(x) : (\{0, 1..., M\} \times \mathbb{N} \times \{0, 1..., M\})^{n-1} \mapsto \{0, 1, ..., M\}$ maps the single-category history, $\{U_{\ell x}, L_{\ell x}, Y_{\ell x} : \ell \le n-1\}$, into actions U_{nx} for category x.

We can convert each sub-problem from a problem with a random finite time horizon to one with an infinite time horizon as follows. First, N follows a geometric distribution with parameter $1-\gamma$, where $\gamma = \frac{s}{s+r}$. Then,

$$E^{\pi(x)}\left[\sum_{n=1}^{N}(Y_{nx} - cZ_{nx})\right] = \gamma E^{\pi(x)}\left[\sum_{n=1}^{\infty}\gamma^{n-1}(Y_{nx} - cZ_{nx})\right].$$

We now solve this MDP using stochastic dynamic programming. We define the value function

$$V_x(\alpha,\beta) = \sup_{\pi(x)\in\Pi(x)} E^{\pi(x)}\left[\sum_{n=1}^{\infty}\gamma^{n-1}(Y_{nx} - cZ_{nx})\Big|\alpha_{0x} = \alpha, \beta_{0x} = \beta\right]. \tag{2}$$

This is a two-dimensional dynamic problem, and we can write its Bellman equation as:

$$V_x(\alpha,\beta) = \max_{u\in\{0,1,...,M\}} Q(\alpha,\beta,u), \tag{3}$$

where $Q(\alpha,\beta,u)$ is the expected reward when we forward at most u items and behave optimally afterwards:

$$\begin{aligned} Q(\alpha,\beta,u) =& E[Y_{1x} - c\min(u, L_{1x}) + \gamma V_x(\alpha_{1x},\beta_{1x})|U_{1x} = u, \alpha_{0x} = \alpha, \beta_{0x} = \beta] \\ =& (\alpha/(\alpha+\beta) - c)E[\min(u, L_{1x})] + \gamma E[V_x(\alpha_{1x},\beta_{1x})|U_{1x} = u, \alpha_{0x} = \alpha, \beta_{0x} = \beta]. \end{aligned} \tag{4}$$

The first in equation (4) is the immediate reward for the action. To compute the second term, we use that L_{nx} is geometric with parameter $\xi_x = \frac{s}{\lambda_x+s}$, and so $E[\min(u, L_{1x})] = \frac{1-\xi_x}{\xi_x}[1 - (1-\xi_x)^u]$.

The third term in equation (4) specifies the expected future reward for forwarding at most u items,

$$\begin{aligned} & E[V_x(\alpha_{1x}, \beta_{1x})|U_{1x} = u, \alpha_{0x} = \alpha, \beta_{0x} = \beta] \\ = & \sum_{i=0}^{u} \sum_{k=0}^{i} P\left(Y_{1x} = k | Z_{1x} = i, \alpha_{0x} = \alpha, \beta_{0x} = \beta\right) P(Z_{1x} = i | U_{1x} = u) V_x(\alpha + k, \beta + i - k). \end{aligned}$$

To compute this, one can show that $P\left(Y_{1x} = k \mid Z_{1x} = i, \alpha_{0x} = \alpha, \beta_{0x} = \beta\right) = \binom{i}{k} \frac{B(\alpha+k, \beta+i-k)}{B(\alpha,\beta)}$, where $B(x, y)$ is the beta function. The distribution $Z_{1x} = \min(U_{1x}, L_{1x})|U_{1x}$ is given by:

$$P(Z_{1x} = i | U_{1x} = u) = \begin{cases} (1 - \xi_x)^i \xi_x & \text{if } 0 \leq i < u, \\ (1 - \xi_x)^u & \text{if } i = u. \end{cases}$$

With these expressions, we can use backward induction, illustrated in Algorithm 1, to solve the dynamic program in equation (2) by considering a truncated time-horizon problem terminated at $\tilde{T}$, and averaging an upper bound on the value function, $V_x^L(\alpha, \beta; \tilde{T})$ and a lower bound, $V_x^U(\alpha, \beta; \tilde{T})$.

3 Mathematical Model with a Constraint on Items Forwarded

In Section 2, we assumed that the system knows the unit cost, c, that the user incurs for reviewing each item. In reality, we often do not know this cost. In this section, we instead assume that the number of items forwarded in each step is constrained, $\sum_{x=1}^{k} U_{nx} \leq M \ \forall n$. Our objective is to maximize the expected number of relevant items forwarded, subject to this constraint:

$$\sup_{\tilde{\pi} \in \tilde{\Pi}} E^{\tilde{\pi}} \left[\sum_{n=1}^{N} \sum_{x=1}^{k} Y_{nx} \right], \tag{5}$$

where $\tilde{\Pi} = \{\pi \in \Pi : \sum_{x=1}^{k} U_{nx} \leq M \ \forall n\}$ and $\tilde{\pi}$ is a policy in $\tilde{\Pi}$ that satisfies the constraint.

In contrast with the previous problem, computation in this problem scales exponentially in k due to "curse of dimensionality", because we can no longer decompose equation (5) into multiple tractable sub-problems (Powell 2007). Instead, we consider a Lagrangian relaxation, following developments Hu et al. (2014), Xie and Frazier (2013b), that provides a computationally tractable upper bound on the value of equation (5), and which motivates an index-based heuristic policy below in Section 3.1.

Let $\boldsymbol{\alpha}_0 = (\alpha_{01}, ..., \alpha_{0k})$ and $\boldsymbol{\beta}_0 = (\beta_{01}, ..., \beta_{0k})$. Let $\boldsymbol{\nu} = (\nu_1, ..., \nu_N)$ be a vector of Lagrange multipliers, with each $\nu_n \geq 0$ denoting a unit cost (or penalty) when we violate the constraint, $\sum_{x=1}^{k} Z_{nx} \leq M$, at step n. We can then write the Lagrangian relaxation of (5) as

$$\begin{aligned} V^{\boldsymbol{\nu}}(\boldsymbol{\alpha}_0, \boldsymbol{\beta}_0) &= \sup_{\pi \in \Pi} E^{\pi} \left[\sum_{n=1}^{N} \left[\sum_{x=1}^{k} Y_{nx} - \nu_n \left(\sum_{x=1}^{k} Z_{nx} - M \right) \right] \right] & (6) \\ &= \sum_{x=1}^{k} \sup_{\pi(x) \in \Pi(x)} E^{\pi(x)} \left[\sum_{n=1}^{N} (Y_{nx} - \nu_n Z_{nx}) \right] + ME \left[\sum_{n=1}^{N} \nu_n \right]. & (7) \end{aligned}$$

Because of the constraint $\sum_{x=1}^{k} Z_{nx} \leq M$ for all n satisfied by $\tilde{\pi} \in \tilde{\Pi}$, and the non-negativity of the Lagrange multiplier, $\boldsymbol{\nu} \geq \mathbf{0}$, the value in equation (6) provides an upper bound on the value of equation (5). Given $\boldsymbol{\nu} \geq \mathbf{0}$, let $V_x^{\boldsymbol{\nu}}(\alpha_{0x}, \beta_{0x}) = \sup_{\pi \in \Pi} E^{\pi}\left[\sum_{n=1}^{N}(Y_{nx} - \nu_n Z_{nx})\right]$. Then $V^{\boldsymbol{\nu}}(\boldsymbol{\alpha}_0, \boldsymbol{\beta}_0)$ can be decomposed into a sum of multiple sub-problems, plus a constant,

$$V^{\boldsymbol{\nu}}(\boldsymbol{\alpha}_0, \boldsymbol{\beta}_0) \;=\; \sum_{x=1}^{k} V_x^{\boldsymbol{\nu}}(\alpha_{0x}, \beta_{0x}) + ME\left[\sum_{n=1}^{N} \nu_n\right] \geq \sup_{\tilde{\pi} \in \tilde{\Pi}} E^{\tilde{\pi}}\left[\sum_{n=1}^{N}\sum_{x=1}^{k} Y_{nx}\right]. \tag{8}$$

This upper bound $V^{\boldsymbol{\nu}}(\boldsymbol{\alpha}_0, \boldsymbol{\beta}_0)$ is useful because it can be computed efficiently, as the sum of independent and easy-to-solve stochastic control sub-problems. Furthermore, if we consider this upper bound for the special case $\boldsymbol{\nu} = \nu \boldsymbol{e}$ with $\nu \geq 0$ and $\boldsymbol{e} = (1, 1, ..., 1)$, then $V_x^{\boldsymbol{\nu}}(\alpha_{0x}, \beta_{0x})$ recovers the total expected reward in the information filtering problem described in Section 2, with a fixed unit cost, ν, for forwarding each item. $V_x^{\boldsymbol{\nu}}(\alpha, \beta)$ can be computed efficiently using Algorithm 1.

We can construct tighter upper bound by taking the infimum of $V^{\boldsymbol{\nu}}(\boldsymbol{\alpha}_0, \boldsymbol{\beta}_0)$ over sets of potential values for our Lagrange multipliers,

$$\mathrm{UB}(\boldsymbol{\alpha}_0, \boldsymbol{\beta}_0) \;=\; \inf_{\boldsymbol{\nu} \geq 0, \boldsymbol{\nu} = \nu \boldsymbol{e}} V^{\boldsymbol{\nu}}(\boldsymbol{\alpha}_0, \boldsymbol{\beta}_0) \geq \inf_{\boldsymbol{\nu} \geq 0} V^{\boldsymbol{\nu}}(\boldsymbol{\alpha}_0, \boldsymbol{\beta}_0). \tag{9}$$

Inequality (9) holds because $\{\boldsymbol{\nu} \geq \mathbf{0} : \boldsymbol{\nu} = \nu \boldsymbol{e}\} \subseteq \{\boldsymbol{\nu} \geq \mathbf{0}\}$.

We conjecture that the function $V_x^{\boldsymbol{\nu}}(\alpha_{0x}, \beta_{0x})$ is convex and non-increasing in $\boldsymbol{\nu}$. The other term, $ME\left[\sum_{n=1}^{N} \nu_n\right]$ is also a convex function of $\boldsymbol{\nu}$, so $V^{\boldsymbol{\nu}}(\boldsymbol{\alpha}_0, \boldsymbol{\beta}_0)$ is conjectured to be convex. If this conjecture is true, we can use a bisection algorithm to find $\boldsymbol{\nu} \geq \mathbf{0}$ in the region $\{\boldsymbol{\nu} \geq \mathbf{0} : \boldsymbol{\nu} = \nu \boldsymbol{e}\}$ that minimizes $V^{\boldsymbol{\nu}}(\boldsymbol{\alpha}_0, \boldsymbol{\beta}_0)$, obtaining $\mathrm{UB}(\boldsymbol{\alpha}_0, \boldsymbol{\beta}_0)$. The conjectured non-increasing property of $V_x^{\boldsymbol{\nu}}(\alpha_{0x}, \beta_{0x})$ induces another property of the optimal solution: given $\nu > 0$, if it is optimal to forward m items at $(\alpha_{0x}, \beta_{0x})$, then it should be optimal to forward at least m items for all $0 \leq \nu' < \nu$.

Given $\boldsymbol{\nu} = \nu \boldsymbol{e}$, let $\pi_{\nu}^{*}(x)$ be the optimal policy for $V_x^{\boldsymbol{\nu}}(\cdot, \cdot)$, and $U^{*}(\alpha, \beta) \in \pi_{\nu}^{*}(x)$ be the optimal decision at state (α, β). For each $0 \leq u \leq M$, we define $\nu^{*}(u, \alpha, \beta)$ to be the largest reward achieved for forwarding (or the largest cost a user would pay to view) at least u number of items at state (α, β),

$$\nu^{*}(u, \alpha, \beta) \;=\; \sup_{\nu \geq 0}\{\nu : U_{\nu}^{*}(\alpha, \beta) \geq u \text{ with } U_{\nu}^{*}(\alpha, \beta) \in \pi_{\nu}^{*}(x)\} \tag{10}$$

$$\;=\; \sup_{\nu \geq 0}\left\{\nu : V^{\boldsymbol{\nu}}(\alpha, \beta) = \max_{u' \geq u} Q^{\nu}(\alpha, \beta, u')\right\}. \tag{11}$$

In the special case $M = 1$ and $L_{nx} \geq 1$ for all n and x, the index $\nu^{*}(1, \alpha, \beta)$ actually corresponds to Gittin's index for a two-armed bandit (Gittins and Gittins 1979, Whittle 1980), where one arm gives the known value of not forwarding (0), and the other arm gives the unknown value of forwarding in a single-category problem. Figure 2 plots the optimal decision $U^{*}(\alpha, \beta)$ against Lagrange multiplier $\nu = \nu \mathbf{e}$ for different values of α, β, ξ, and γ. In both plots, we see that $U^{*}(\alpha, \beta)$ is non-increasing

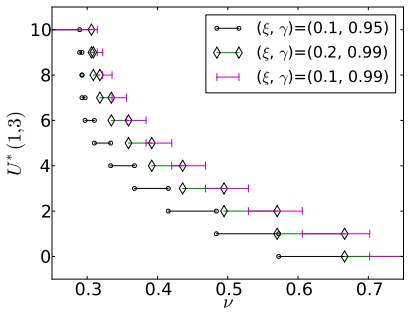

(a) Plot of $U^*(\alpha, \beta)$ vs. ν at state $(\alpha, \beta) = (1, 3)$ for different values of ξ and γ.

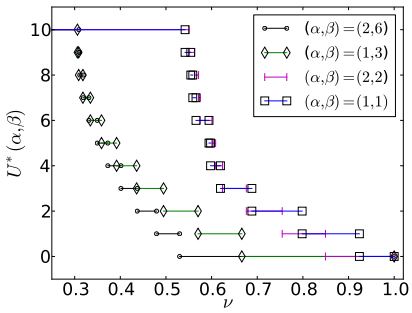

(b) Plot of $U^*(\alpha, \beta)$ vs. ν at $\xi = 0.2$ and $\gamma = 0.99$ for different pairs of (α, β).

Figure 2: Plots of optimal decision $U^*(\alpha, \beta)$ against Lagrange multiplier $\boldsymbol{\nu} = \nu \boldsymbol{e}$ for various values of $(\alpha, \beta, \xi, \gamma)$ with $M = 10$. In each case, we solve $V_x^{\nu}(\alpha, \beta)$ via Algorithm 1 to find $U^*(\alpha, \beta)$ for each $\nu \in [0, 1]$. Here, $\nu^*(u, \alpha, \beta)$ is ν at the right-most endpoint of the interval where $U^*(\alpha, \beta) \geq u$.

in ν. The observation confirms the conjectured statement that if it is optimal to forward m items at ν, then it is optimal to forward at least m items for all $\nu' < \nu$. Furthermore, $\nu^*(u, \alpha, \beta)$ is ν at the right-most endpoint of intervals where $U^*(\alpha, \beta) \geq u$. In Figure 2(a), we vary ξ and γ while fixing (α, β). Increasing the discount factor, γ, increases $\nu^*(u, \alpha, \beta)$ because a user spends longer in the system. In contract, as ξ increases, $\nu^*(u, \alpha, \beta)$ becomes smaller because on average there are fewer items queued in the system. Figure 2(b) illustrates that $\nu^*(u, \alpha', \beta') \geq \nu^*(u, \alpha, \beta)$ in the case when $\frac{\alpha'}{\alpha'+\beta'} = \frac{\alpha}{\alpha+\beta}$ but $\alpha' + \beta' < \alpha + \beta$. This reflects that exploration provides the largest value when observation size is small, and uncertainty is high.

3.1 MDP-based Information Filtering (MDP-IF) Policy

In this section, we propose a heuristic index-based policy for ranking items queued at step n, disregarding whether the user cost is known. We call our policy the MDP-based Information Filtering Policy, abbreviated *MDP-IF*. We then discuss scenarios where the proposed policy is Bayes-optimal. At each step n, we first compute $\nu^*(u, \alpha_{nx}, \beta_{nx})$ for all possible $u \leq M$ in each category x, then rank items from the categories based on the computed $\nu^*(u, \alpha_{nx}, \beta_{nx})$, among all categories and all $u \leq M$. $\nu^*(u, \alpha_{nx}, \beta_{nx})$ reflects the highest reward one can achieve if we choose to forward at least u items from the category x at the current state, $(\alpha_{nx}, \beta_{nx})$. In the situation where we have a constraint that $\sum_{x=1}^{k} Z_{nx} \leq M$, we would forward items from the ranked list until we exhaust M slots. This heuristic index policy is exactly the optimal policy corresponding to the Gittins index in a conventional multi-armed bandit problem, when there is at least one item queued in each category per step and $M = 1$. On the other hand, if a user cost c is specified, then we would forward items from the ranked list with $\nu^*(u, \alpha_{nx}, \beta_{nx}) \geq c$. In a general case with both c and M specified, we follow the same procedure with a stop criteria if any condition is violated.

Figure 3 demonstrates how this index policy works in a two-category problem (category "O"

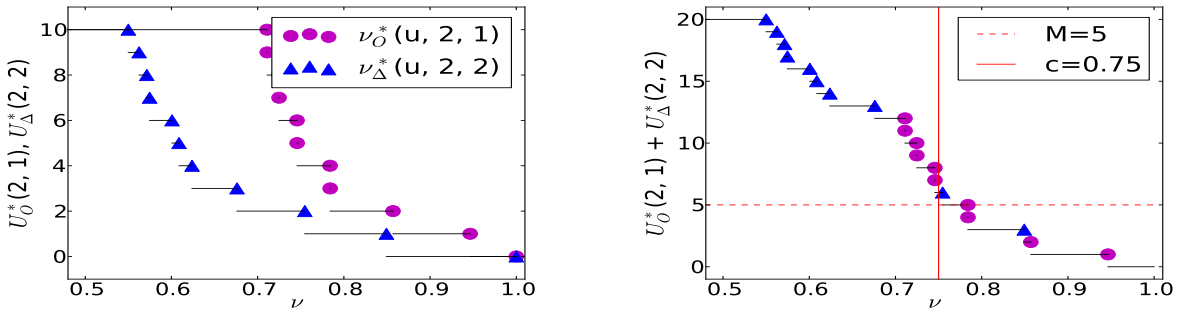

(a) Plot of $U_O^*(2,1)$ and $U_\Delta^*(2,2)$ against ν.

(b) Plot of $U_O^*(2,1) + U_\Delta^*(2,2)$ against ν

Figure 3: There are two categories: category O at state (2, 1) and category Δ at state (2, 2). Figure (a) shows optimal decision $U_O^*(2,1)$, $U_\Delta^*(2,2)$ against lagrangian multiplier ν, and also plots $\nu_O^*(u,2,1)$ (denoted in purple circles) and $\nu_\Delta^*(u,2,2)$ (denoted in blue triangles). Figure (b) plots $U_O^*(2,1) + U_\Delta^*(2,2)$ vs. ν, with ranked list of $\nu_O^*(u,2,1)$ and $\nu_\Delta^*(u,2,2)$ among all $u \in \{0,1,...,10\}$. When $M = 5$, the rank list would be: $\{O, O, \Delta, O, O\}$. When the user cost is 0.75, the rank list would be the six right-most items, $\{O, O, \Delta, O, O, \Delta\}$.

and category "Δ"). Category "O" is at state (2, 1) and category "Δ" is at state (2, 2). Figure (a) plots optimal decisions $U_O^*(2,1)$ and $U_\Delta^*(2,2)$ as a function of ν, and identifies their $\nu_O^*(u,2,1)$ and $\nu_\Delta^*(2,2)$ values. Then, Figure (b) plots $U_O^*(2,1) + U_\Delta^*(2,2)$ against ν, with a ranked list of $\nu_O^*(u,2,1)$ and $\nu_\Delta^*(u,2,2)$ among all $u \in \{0,1,...,10\}$. When $M = 5$, the ranked list would be items below the dashed line: $\{O, O, \Delta, O, O\}$. When the user cost is 0.75, the ranked list would be the six items on the right of the vertical line (defined by $c = 0.75$): $\{O, O, \Delta, O, O, \Delta\}$.

4 Experimental Results

In this section, we show numerical results of Monte Carlo simulations in four different parameter settings to illustrate how this heuristic index-based policy performs compared to other competing methods, including the pure exploitation and upper confidence bound (UCB) policies. At each step $n+1$, pure exploitation ranks items based on posterior means of user preference across categories, while UCB ranks items by $(1 - 1/t_n)$ % quantiles from the posterior distribution on user preferences, with t_n denoting the total number of items shown by n.

We run experiments for four different scenarios, each with a chosen set of parameters, including γ and $\{\alpha_{0x}, \beta_{0x}, \xi_x\}_{x \in \{1,...,k\}}$. Each scenarios consists of simulations for a range of $k \in \{5, 10, 30, 50, 100\}$, for us to understand how each policy behaves as category size, k, varies. Specific parameters setting are described in sub-figure captions in Figure 4. Here, γ is chosen based on the estimated parameter of the empirical distribution of user visits in astro-ph.GA and astro-ph.CO from the arXiv.org dataset in 2009-2010. Similarly, ξ_x is chosen based on the empirical distributions. For convenience, α_{0x} and β_{0x} is chosen to be 1 for three scenarios, and 5 for the last scenario to understand how the policy performs when we are more certain about user preferences. For constraints, we set $M = 5$, the maximum number of items forward, and unit cost as $c = 0.49$,

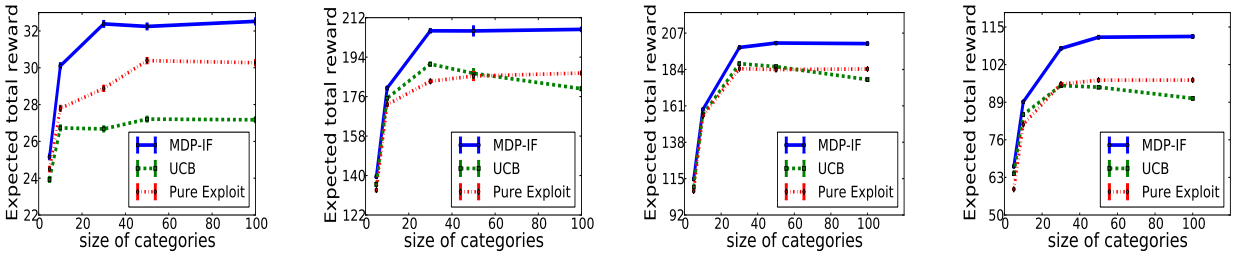

(a) $M = 5$, $c = .49$, $\gamma = \mathbf{0.95}$, $\alpha_{0x} = 1$, $\beta_{0x} = 1$, $\xi_x = 0.1\ \forall x$

(b) $M = 5$, $c = .49$, $\gamma = \mathbf{0.99}$, $\alpha_{0x} = 1$, $\beta_{0x} = 1$, $\xi_x = 0.1\ \forall x$

(c) $M = 5$, $c = .49$, $\gamma = 0.99$, $\alpha_{0x} = 1$, $\beta_{0x} = 1$, $\xi_{\mathbf{x}} = \mathbf{0.2}\ \forall x$

(d) $M = 5$, $c = .49$, $\gamma = 0.99$, $\alpha_{\mathbf{0x}} = \mathbf{5}$, $\beta_{\mathbf{0x}} = \mathbf{5}$, $\xi_x = 0.1\ \forall x$

Figure 4: Each sub-figure shows plots of expected total reward with 95% confidence intervals under each policy (MDP-IF in solid blue line, UCB in red dashed line, Pure Exploit in green dotted line) against category size in a given parameter setting specified in its sub-caption. $50,000$ users are simulated in each scenario.

which is motivated by our conjecture that exploration adds the most benefit when cost is near the initial mean of θ_x, which is 0.5 in our test cases.

With the chosen set of parameters in each scenario, we simulate a large set of users, with user visits, items arrivals between two visits, and user feedbacks on forwarded items generated according to the model formulated in Section 2. At each visit, each policy (MDP-IF, UCB, or pure exploitation) decides how many and which items to forward. The reward is then collected to compute the expected total reward for each policy.

Each line in Figure 4 shows the total expected reward with 95% confidence intervals against category size, k, for a given parameter setting. In all scenarios, the MDP-IF policy outperforms both UCB and pure exploitation, with magnified improvement for a larger category size and, mostly, a non-decreasing relationship with k. For some larger k, the improvement from the MDP-IF policy widens as γ increases (comparing Figure 4(a) and (b)) or ξ_x shrinks (comparing Figure 4(b) and (c)), or the prior user preference has less variance (comparing Figure 4(b) and (d)).

5 Conclusion

We consider the information filtering problem where we face a voluminous stream of items and need to decide sequentially on which batch of items to forward to a user so that the total reward, the number of relevant items shown minus the cost of forwarded items, is maximized. With a focus on limited historical data, we formulate the problem as a Markov Decision Process in a Bayesian setting, and then provide a computationally tractable algorithm that is Bayes-optimal. In a setting where the total number of item shown is constrained, we consider a Lagrangian relaxation of the problem and provide an index-based policy that ranks items. As shown in the numerical section, our index-based policy outperforms UCB and the pure exploitation policy, and provides magnified benefits in many settings.

Acknowledgements

Authors were supported by NSF IIS-1247696 and IIS-1247696. Peter Frazier was also supported by NSF CAREER CMMI-1254298, AFOSR YIP FA9550-11-1-0083, and AFOSR FA9550-12-1-0200.

References

Agarwal, D., B.C. Chen, P. Elango. 2009. Explore/exploit schemes for web content optimization. *Proceedings of 2009 Ninth IEEE International Conference on Data Mining*. ICDM '09, IEEE Computer Society, Washington, DC, USA, 1–10.

Araman, V.F., R. Caldenty. 2009. Dynamic pricing for perishable products with demand learning. *Operations Research* **57**(5) 1169 – 1188.

arXiv.org. 2014. arxiv.org e-print archive. `http://arxiv.org`. Accessed: 2014-06-30.

Auer, P., N. Cesa-Bianchi, P. Fischer. 2002. Finite-time analysis of the multiarmed bandit problem. *Machine learning* **47**(2) 235–256.

Auer, P., N. Cesa-Bianchi, Y. Freund, RE Schapire. 1995. Gambling in a rigged casino: The adversarial multi-armed banditproblem. *Foundations of Computer Science, 1995. Proceedings., 36th Annual Symposium on*. 322–331.

Besbes, O., A. Zeevi. 2009. Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Operations Research* **57**(6) 1407–1420.

den Boer, A.V., B. Zwart. 2013. Simultaneously learning and optimizing using controlled variance pricing. *Management Science* **60**(3) 770–783.

Ding, X., M.L. Puterman, A. Bisi. 2002. The Censored Newsvendor and the Optimal Acquisition of Information. *Operations Research* **50**(3) 517–527.

Frazier, P. I., W. B. Powell, S. Dayanik. 2008. A knowledge gradient policy for sequential information collection. *SIAM Journal on Control and Optimization* **47**(5) 2410–2439.

Gittins, Author(s) J. C., J. C. Gittins. 1979. Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society, Series B* 148–177.

Gittins, J., K. Glazebrook, R. Weber. 2011. *Multi-armed Bandit Allocation Indices*. 2nd ed. Wiley.

Gittins, J. C., D. M. Jones. 1974. A dynamic allocation index for the sequential design of experiments. J. Gani, ed., *Progress in Statistics*. North-Holland, Amsterdam, 241–266.

Hofmann, K., S. Whiteson, M. Rijke. 2013. Balancing exploration and exploitation in listwise and pairwise online learning to rank for information retrieval. *Information Retrieval* **16**(1) 63–90.

Hu, W., P. I. Frazier, J. Xie. 2014. Parallel bayesian policies for finite-horizon multiple comparisons with a known standard. A. Tolk, S. D. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, J. A. Miller, eds., *Proceedings of the 2014 Winter Simulation Conference*. Institute of Electrical and Electronics Engineers, Inc., Piscataway, New Jersey.

Jaksch, T., R. Ortner, P. Auer. 2010. Near-optimal regret bounds for reinforcement learning. *Journal of Machine Learning Research (JMLR)* **11** 1563–1600.

Kaelbling, L.P., M.L. Littman, A.R. Cassandra. 1998. Planning and acting in partially observable stochastic domains. *Artificial Intelligence* **101**(1-2) 99–134.

Lariviere, M.A., E.L. Porteus. 1999. Stalking Information: Bayesian Inventory Management with Unobserved Lost Sales. *Management Science* **45**(3) 346–363.

Powell, W.B. 2007. *Approximate Dynamic Programming: Solving the Curses of Dimensionality (Wiley Series in Probability and Statistics)*. Wiley-Interscience, Hoboken, NJ, USA.

Powell, W.B., A. Ruszczyn'ski, H. Topaloglu. 2004. Learning algorithms for separable approximations of discrete stochastic optimization problems. *Mathematics of Operations Research* **29**(4) 814–836.

Sutton, R.S., A.G. Barto. 1998. *Introduction to Reinforcement Learning*. MIT Press, Cambridge, MA, USA.

Whittle, P. 1980. Multi-armed bandits and the Gittins index. *Journal of the Royal Statistical Society. Series B (Methodological)* **42**(2) 143–149.

Xie, J., P. I. Frazier. 2013a. Sequential bayes-optimal policies for multiple comparisons with a known standard. *Operations Research* **61**(5) 1174–1189. doi:10.1287/opre.2013.1207.

Xie, J., P. I. Frazier. 2013b. Upper bounds on the bayes-optimal procedure for ranking & selection with independent normal priors. R. Pasupathy, S.-H. Kim, A. Tolk, R. Hill, M. E. Kuhl, eds., *Proceedings of the 2013 Winter Simulation Conference*. Institute of Electrical and Electronics Engineers, Inc., Piscataway, New Jersey, 877–887.

Yue, Y., J. Broder, R. Kleinberg, T. Joachims. 2009. The K-armed Dueling Bandits Problem. *Conference on Learning Theory (COLT)*.

Zhang, Y., W. Xu, J. Callan. 2003. Exploration and exploitation in adaptive filtering based on bayesian active learning. *Proceedings of the 20th International Conference*. ICML '03, ACM, Washington, DC, USA, 896–903.

Zhao, X., P. Frazier. 2014. Exploration vs. exploitation in the information filtering problem. *arxiv.org/pdf/1407.8186* .